CLAIMS

1. (currently amended) A method, comprising:

receiving one or more demands for service in a mesh network comprising a plurality of nodes interconnected by a plurality of links;

specifying a threshold corresponding to a maximum number of failure-related cross-connections at a node in the network; and

mapping each of the one or more demands onto a primary path and a restoration path in the network to generate a path plan for the one or more demands in the network, wherein:

reduction of a portion of restoration time associated with failure-related cross-connections in the network is taken into account during the mapping; and

the mapping generates the path plan based on the specified threshold such that, for all nodes in the mesh network, the number of failure-related cross-connections at each node is ~~less~~ no more than the specified threshold.

2. (canceled)

3. (previously presented) The method of claim 1, wherein the mapping results in a maximum number of failure-related cross-connections at all nodes in the network being within a specified tolerance of a theoretical minimum.

4. (previously presented) The method of claim 3, wherein a graph-theoretic condition is used to derive the theoretical minimum.

5. (previously presented) The method of claim 4, wherein the theoretical minimum is defined by $\max_{n \in N} \{\lceil \delta_n / d_n \rceil\}$ where n, a node in the network, is an element of N, the set of all nodes in the network, δ_n is the number of unit demands terminated on node n, and d_n is the number of edges incident on node n.

6. (previously presented) The method of claim 1, wherein the mapping sequentially evaluates each possible path plan for each of the one or more demands and selects the path plan having a smallest maximum number of failure-related cross-connections.

7. (previously presented) The method of claim 1, wherein the mapping comprises:

selecting two node-disjoint paths for each demand, wherein leveling of link loads is taken into account during the selecting; and

for each demand, identifying one of the two node-disjoint paths as the primary path and the other as the restoration path, wherein a maximum number of failure-related cross-connections at all nodes in the network is taken into account during the identifying.

8. (previously presented) The method of claim 7, wherein:

selecting the two node-disjoint paths for each demand minimizes maximum link bandwidth in the network; and

identifying the primary and restoration paths for each demand results in the maximum number of failure-related cross-connections at all nodes in the network being within a specified tolerance of a theoretical minimum.

9. (previously presented) The method of claim 8, wherein a tent pole condition is used to derive the theoretical minimum.

10. (previously presented) The method of claim 7, wherein the selecting of the two node-disjoint paths for each demand and the identifying, for each demand, of the one of the two node-disjoint paths as the primary path and the other as the restoration path are implemented using mixed-integer programming.

11. (previously presented) The method of claim 7, wherein the selecting of the two node-disjoint paths for each demand and the identifying, for each demand, of the one of the two node-disjoint paths as the primary path and the other as the restoration path are implemented using genetic programming.

12. (previously presented) The method of claim 7, wherein the selecting of the two node-disjoint paths for each demand and the identifying, for each demand, of the one of the two node-disjoint paths as the primary path and the other as the restoration path are implemented using a commercial solver.

13. (previously presented) The method of claim 1, wherein the mapping involves demand bundling, wherein demands having a common source node and a common destination node are grouped and routed along a single pair of disjoint primary and restoration paths and at least a portion of connection signaling for the group is carried out jointly.

14. (previously presented) The method of claim 1, wherein the mapping involves traffic aggregation, wherein multiple low-rate channels in the network are consolidated into a high-rate channel and rerouting of the high-rate channel requires fewer cross-connections than rerouting of the multiple low-rate channels.

15. (currently amended) A network manager for a mesh network comprising a plurality of nodes interconnected by a plurality of links, the network manager comprising:

means for receiving one or more demands for service in the network;

means for specifying a threshold corresponding to a maximum number of failure-related cross-connections at a node in the network; and

means for mapping each of the one or more demands onto a primary path and a restoration path in the network to generate a path plan for the one or more demands in the network, wherein:

reduction of a portion of restoration time associated with failure-related cross-connections in the network is taken into account during the mapping; and

the means for mapping generates the path plan based on the specified threshold such that, for all nodes in the mesh network, the number of failure-related cross-connections at each node is ~~less~~ no more than the specified threshold.

16. (canceled)

17. (previously presented) The network manager of claim 15, wherein the path plan results in a maximum number of failure-related cross-connections at all nodes in the network being within a specified tolerance of a theoretical minimum.

18. (previously presented) The network manager of claim 17, wherein a graph-theoretic condition is used to derive the theoretical minimum.

19. (previously presented) The network manager of claim 18, wherein the theoretical minimum is defined by: $\max_{n \in N} \{\lceil \delta_n / d_n \rceil\}$ where n, a node in the network, is an element of N, the set of all nodes in the network, δ_n is the number of unit demands terminated on node n, and d_n is the number of edges incident on node n.

20. (previously presented) The network manager of claim 15, wherein the network manager comprises means for sequentially evaluating each possible path plan for each of the one or more demands and means for selecting the path plan having a smallest maximum number of failure-related cross-connections.

21. (previously presented) The network manager of claim 15, wherein the network manager comprises:

means for performing selection of two node-disjoint paths for each demand, wherein leveling of link loads is taken into account during the selection; and

means for identifying, for each demand, one of the two node-disjoint paths as the primary path and the other as the restoration path, wherein a maximum number of failure-related cross-connections at all nodes in the network is taken into account during the identifying.

22. (previously presented) The network manager of claim 21, wherein:

the means for performing the selection of the two node-disjoint paths for each demand minimizes maximum link bandwidth in the network; and

the means for identifying the primary and restoration paths for each demand results in the maximum number of failure-related cross-connections at all nodes in the network being within a specified tolerance of a theoretical minimum.

23. (previously presented) The network manager of claim 22, wherein a tent pole condition is used to derive the theoretical minimum.

24. (previously presented) The network manager of claim 21, wherein the means for performing the selection and the means for identifying the primary and restoration paths are implemented using mixed-integer programming.

25. (previously presented) The network manager of claim 21, wherein the means for performing the selection and the means for identifying the primary and restoration paths are implemented using genetic programming.

26. (previously presented) The network manager of claim 21, wherein the means for performing the selection and the means for identifying the primary and restoration paths are implemented using a commercial solver.

27. (previously presented) The network manager of claim 15, wherein the network manager comprises means for considering demand bundling in the generation of the path plan, wherein demands having a common source node and a common destination node are grouped and routed along a single pair of disjoint primary and restoration paths and at least a portion of connection signaling for the group is carried out jointly.

28. (previously presented) The network manager of claim 15, wherein the network manager comprises means for considering traffic aggregation in the generation of the path plan, wherein multiple low-rate channels in the network are consolidated into a high-rate channel and rerouting of the high-rate channel requires fewer cross-connections that rerouting of the multiple low-rate channels.

29. (previously presented) A method, comprising:

receiving one or more demands for service in a mesh network comprising a plurality of nodes interconnected by a plurality of links; and

mapping each of the one or more demands onto a primary path and a restoration path in the network to generate a path plan for the one or more demands in the network, wherein:

reduction of a portion of restoration time associated with failure-related cross-connections in the network is taken into account during the mapping;

the mapping results in a maximum number of failure-related cross-connections at all nodes in the network being within a specified tolerance of a theoretical minimum;

a graph-theoretic condition is used to derive the theoretical minimum; and

the theoretical minimum is defined by $\max_{n \in N} \{\lceil \delta_n / d_n \rceil\}$ where n, a node in the network, is an element of N, the set of all nodes in the network, δ_n is the number of unit demands terminated on node n, and d_n is the number of edges incident on node n.

30. (previously presented) A network manager for a mesh network comprising a plurality of nodes interconnected by a plurality of links, the network manager comprising:

means for receiving one or more demands for service in the network; and

means for mapping each of the one or more demands onto a primary path and a restoration path in the network to generate a path plan for the one or more demands in the network, wherein:

reduction of a portion of restoration time associated with failure-related cross-connections in the network is taken into account during the mapping;

the path plan results in a maximum number of failure-related cross-connections at all nodes in the network being within a specified tolerance of a theoretical minimum;

a graph-theoretic condition is used to derive the theoretical minimum; and

the theoretical minimum is defined by: $\max_{n \in N} \{\lceil \delta_n / d_n \rceil\}$ where n, a node in the network, is an element of N, the set of all nodes in the network, δ_n is the number of unit demands terminated on node n, and d_n is the number of edges incident on node n.